

UNITED STATES PATENT APPLICATION

for

A MECHANISM FOR MANAGING POWER GENERATED IN A COMPUTER SYSTEM

Inventors:

James Hermerding

Randy Stanley

Prepared by:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

12400 Wilshire Boulevard

Los Angeles, CA 90025-1026

(303) 740-1980

File No.: 042390.P9249

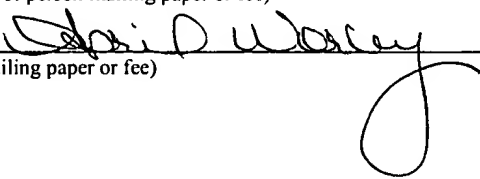
"Express Mail" mailing label number EL750127652US

Date of Deposit December 29, 2000

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231.

April M. Worley

(Typed or printed name of person mailing paper or fee)



(Signature of person mailing paper or fee)

A MECHANISM FOR MANAGING POWER GENERATED IN A COMPUTER SYSTEM

FIELD OF THE INVENTION

The present invention relates to computer systems; more particularly, the present invention relates to power management of computer systems.

BACKGROUND

5 Traditionally, the power generated by microprocessors in mobile computers systems (e.g., notebook computers) was of little concern because of the relatively low speeds at which they operate. However, with the continuous increase of the operating speeds of microprocessors, the power generated by the microprocessor makes cooling the computer system more difficult. For example, typical microprocessors in mobile
10 computer systems generate between 20 – 30 watts in a one-inch form factor. The generation of this magnitude of power at a small location may potentially create thermal issues at the memory device.

Figure 2 illustrates an exemplary cooling system used in notebook computers. The cooling system includes a block coupled to a microprocessor, a heat pipe, a heat
15 exchanger and a cooling fan. Heat generated by the microprocessor is distributed to the heat pipe, which in turn, transfers the heat to the heat exchanger. Subsequently, the heat exchanger is cooled by air blown by the cooling fan. The problem with conventional cooling systems is that it is difficult to dissipate heat generated by more powerful microprocessors in such a small area. Therefore, a method and apparatus for managing
20 the power generated by microprocessors is desired.

5

10

10

10

DETAILED DESCRIPTION

A method and apparatus for managing power generated by microprocessors is described. In the following detailed description of the present invention numerous specific details are set forth in order to provide a thorough understanding of the present invention. However, it will be apparent to one skilled in the art that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form, rather than in detail, in order to avoid obscuring the present invention.

Reference in the specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the invention. The appearances of the phrase “in one embodiment” in various places in the specification are not necessarily all referring to the same embodiment.

Figure 1 is a block diagram of one embodiment of a computer system 100.

According to one embodiment, computer system 100 is a mobile computer system (e.g., laptop, notebook, etc.). Computer system 100 includes central processing units (processors) 105a – 105d coupled to a processor bus 110. In one embodiment, processors 105 are processors in the Pentium® family of processors including the Pentium® II family and mobile Pentium® and Pentium® II processors available from Intel Corporation of Santa Clara, California. Alternatively, other processors may be used.

Chip set 120 is also coupled to processor bus 110. Chip set 120 may include a memory controller for controlling a main memory 113. Further, chipset 120 may also include an Accelerated Graphics Port (AGP) Specification Revision 2.0 interface

developed by Intel Corporation of Santa Clara, California. Chip set 120 is coupled to a video device 125 and handles video data requests to access main memory 113. One of ordinary skill in the art will appreciate that, in other embodiments, each processor 105 may be directly coupled to chipset 120, rather than via processor bus 110.

5 Main memory 113 is coupled to processor bus 110 through chip set 120. Main memory 113 stores sequences of instructions that are executed by processor 105. In one embodiment, main memory 113 includes a dynamic random access memory (DRAM) system; however, main memory 113 may have other configurations. The sequences of instructions executed by processor 105 may be retrieved from main memory 113 or any
10 other storage device. Additional devices may also be coupled to processor bus 110, such as multiple processors and/or multiple main memory devices. Video device 125 is also coupled to chip set 120. In one embodiment, video device includes a video monitor such as a cathode ray tube (CRT) or liquid crystal display (LCD) and necessary support circuitry.

15 Processor bus 110 is coupled to system bus 130 by chip set 120. In one embodiment, system bus 130 is a Peripheral Component Interconnect (PCI) Specification Revision 2.1 standard bus developed by Intel Corporation of Santa Clara, California; however, other bus standards may also be used. Multiple devices, such as audio device 127, may be coupled to system bus 130.

20 Bus bridge 140 couples system bus 130 to secondary bus 150. In one embodiment, secondary bus 150 is an Industry Standard Architecture (ISA) Specification Revision 1.0a bus developed by International Business Machines of Armonk, New York. However, other bus standards may also be used, for example Extended Industry Standard

Architecture (EISA) Specification Revision 3.12 developed by Compaq Computer, et al. Multiple devices, such as hard disk 153 and disk drive 154 may be coupled to secondary bus 150. Other devices, such as cursor control devices (not shown in **Figure 1**), may be coupled to secondary bus 150.

5 **Figure 3** illustrates one embodiment of a cooling system 300 within computer system 100. Cooling system 300 includes blocks 310, heat pipe 320, heat exchanger 330 and cooling fan 340. A block 310 is coupled to each of the processors 105. According to one embodiment, blocks 310 are made of copper. However, one of ordinary skill in the art will appreciate that blocks 310 may be made of other materials. Heat pipe 320 is
10 coupled to each processor 105 via blocks 310.

 According to one embodiment, heat pipe 320 is a hollow copper tube filled with a small amount of liquid such as water. In a further embodiment, heat pipe 320 maintains a vacuum. Since water boils rapidly in a vacuum, the water becomes vapor upon being heated by a processor 105, and is transferred away from the point where the heat is being
15 generated. Therefore, heat generated by each processor 105 is transferred by heat pipe 320. Heat exchanger 330 dissipates the heat transferred by heat pipe 320. Cooling fan 340 further dissipates the heat by blowing air across heat exchanger 330.

 According to one embodiment, computer system 100 is arranged such that instruction tasks are moved between processor 105a – 105d based upon the heat being
20 generated at each. In such an embodiment, each processor 105 includes a thermal sensor that provides thermal feedback to the operating system that runs on computer system 100. Based upon the feedback, the operating system makes decisions on how to partition the workload among the processors 105.

Figure 4 is a flow diagram of one embodiment for the operation of computer system 100. At process block 410, the operating system for computer system 100 monitors the currently active processor 105 to determine the thermal state. In one embodiment, the operating system receives a thermal signal from the active processor 105 once the processor 105 has reached $\frac{1}{4}$ of its power capacity. However, in other embodiments the thermal signal may be transmitted upon reaching other increment levels of the power capacity of a processor 105.

At process block 420, it is determined whether the active processor 105 is generating the thermal signal. If the thermal signal is not being transmitted, the active processor 105 is operating below the predetermined thermal threshold. As a result, control is returned to process block 410 where the operating system continues to monitor the active processor 105. If, however, it is determined that the thermal signal is being transmitted, the least recently used (LRU) processor 105 in computer system 100 is determined, process block 430.

According to one embodiment, the LRU processor 105 is the processor 105 that has been inactive for the longest interval of time. At process block 440, computer system 100 operations continue at the new active processor 105 (e.g., the LRU processor 105). Moving the processor 105 workload between multiple processors 105 distributes the heat generated by the processors within computer system 100. For example, rather than having one processor 105 generate 20 watts of power, the 20 watts may be distributed evenly between multiple processors 105. In the present embodiment, computer system 100 operates so that 5 watts is generated by each of the processors 105a – 105d. Cooling system 100 may more easily distribute the four different 5-watt sources than one 20-watt

source.

In another embodiment, computer system 100 distributes tasks between processor 105a – 105d based upon the heat being generated at each. In such an embodiment, the operating system includes multiple threads that partition the workload so that one processor 105 does not overheat. Based upon thermal feedback received from each processor 105, the operating system prioritizes the workload based upon the coolest processor 105. Distributing instruction tasks between processors 105 enables cooling system 300 to more easily dissipate heat generated by processors 105.

Whereas many alterations and modifications of the present invention will no doubt become apparent to a person of ordinary skill in the art after having read the foregoing description, it is to be understood that any particular embodiment shown and described by way of illustration is in no way intended to be considered limiting. Therefore, references to details of various embodiments are not intended to limit the scope of the claims which in themselves recite only those features regarded as the invention.

Therefore, a mechanism for managing the power generated by microprocessors has been described.